

The Behavior of Epidemics under Bounded Susceptibility

Subhashini Krishnasamy
Department of ECE
UT Austin
subhashini.kb@utexas.edu

Siddhartha Banerjee
Department of MS&E
Stanford University
sidb@stanford.edu

Sanjay Shakkottai
Department of ECE
UT Austin
shakkott@austin.utexas.edu

ABSTRACT

We investigate the sensitivity of epidemic behavior to a *bounded susceptibility* constraint – susceptible nodes are infected by their neighbors via the regular SI/SIS dynamics, but subject to a *cap on the infection rate*. Such a constraint is motivated by modern social networks, wherein messages are broadcast to all neighbors, but attention spans are limited. Bounded susceptibility also arises in distributed computing applications with download bandwidth constraints, and in human epidemics under quarantine policies.

Network epidemics have been extensively studied in literature; prior work characterizes the graph structures required to ensure fast spreading under the SI dynamics, and long lifetime under the SIS dynamics. In particular, these conditions turn out to be meaningful for two classes of networks of practical relevance – dense, uniform (i.e., *clique-like*) graphs, and sparse, structured (i.e., *star-like*) graphs. We show that bounded susceptibility has a surprising impact on epidemic behavior in these graph families. For the SI dynamics, bounded susceptibility has no effect on star-like networks, but dramatically alters the spreading time in clique-like networks. In contrast, for the SIS dynamics, clique-like networks are unaffected, but star-like networks exhibit a sharp change in extinction times under bounded susceptibility.

Our findings are useful for the design of disease-resistant networks and infrastructure networks. More generally, they show that results for existing epidemic models are sensitive to modeling assumptions in non-intuitive ways, and suggest caution in directly using these as guidelines for real systems.

Keywords

Network Epidemics; SI/SIS Dynamics; Non-linear Dynamics; Bounded Susceptibility

1. INTRODUCTION

Epidemic processes – stochastic models for spreading in graphs – play an important role in a variety of domains ranging from social networks (opinion spread), epidemiology (disease spread) and distributed computing (information spread). Though there are several models in literature for such processes, a characteristic feature of these is the spread of the process from ‘infected’ to ‘non-infected’ nodes along the edges of the network. In this work, we investigate the behavior of such processes under the assumption that *the maximum susceptibility of a node to infection from its neighbors is bounded*. This constraint seems natural (and

we present several settings where it is so) – however we show that imposing it brings to light some surprising behavior of epidemic processes in different networks.

In this paper, we focus on two major classes of epidemic processes: the *SI* process where nodes never recover from infection, and the *SIS* process where nodes recover, but can get reinfected. The *SI* or ‘Susceptible-Infected’ dynamics [1, 2, 3, 4] is the most common model for one-way dissemination. Nodes exist in 2 states – ‘susceptible’ (*S*) and ‘infected’ (*I*). Infected nodes never recover, and try to infect each neighboring susceptible node at a rate β (after a random delay, drawn from an i.i.d. $\text{Exponential}(\beta)$ distribution). We refer to β as the *virulence* of the infection process. The *SI* dynamics terminate when all nodes are infected, and the corresponding time interval is called the *spreading time*. In particular, we are interested in networks with small spreading time (ideally $O(1)$ in the number of nodes).

A small modification of this model gives the popular *SIS* or ‘Susceptible-Infected-Susceptible’ dynamics [5, 6, 7, 8, 9, 10, 11] for epidemic processes. This is identical to the *SI* dynamics, except that now an infected node can revert to being susceptible at a (normalised) rate of ‘1’. The absorbing state for the *SIS* epidemic is when all nodes become susceptible; the time taken to reach this state is called the *extinction time*. We are interested in networks where this extinction time is exponentially large (i.e., $\Omega(e^{\text{poly}(n)})$), where n is the number of nodes in the network).

A key assumption in both models is that *the infection rate for a susceptible node scales linearly in the size of the infected neighborhood*; formally, a susceptible node j transitions to being infected at a rate $\beta|I_j|$, where I_j is the set of infected neighbors of j . This assumption makes the models conceptually simple and tractable. However, it also means that in these models, high-degree nodes may get infected at unrealistically high rates. In several real-life settings, this is unreasonable; for example, in social networks, users broadcast content to all their neighbors, but attention-span constraints limit consumption of content to only a few neighbors. Similarly, in sensor networks using epidemic protocols, nodes can broadcast easily, but have download bandwidth constraints, and so can listen to only some of their neighbors. Refer to Section 1.2 for more examples.

To modify epidemic dynamics to fit such settings, we need to assume that nodes have *bounded susceptibility*, i.e., a maximum rate at which they can get infected, *which is independent of network size*. We can capture this via the notion of an *infection-profile*: a susceptible node j , we assume, is infected at a rate $\Phi(\beta|I_j|)$, where $\Phi(\cdot)$ is a *bounded* function.

Now, in order to use epidemic models to understand real-life processes, one hopes that the behavior of these epidemics is somewhat insensitive to the infection-profile Φ . This motivates the central question we address in this paper:

How robust is SI/SIS epidemic behavior under an additional bounded susceptibility constraint?

The answer turns out to depend on the dynamics and network structure in a surprisingly subtle way. In particular, bounded susceptibility exposes a sharp dichotomy between two classes of networks of practical relevance:

- *Star-like networks*, i.e., sparse, structured graphs, with a small number of edges, mostly concentrated among a dense core, with all other nodes at a short distance from the core. These are good models for *engineered networks*, which are designed to be highly connected while using only a small number of edges.
- *Clique-like networks*, i.e., dense, uniform graphs, with edges spread evenly among nodes. A good example is the connected-regime Erdős-Rényi graph (i.e., $G(n, p)$ with $p = \Omega(\log n/n)$). These are good models for *organically-formed networks*.

In prior work, the distinction between the two classes is suggested, but not explicitly made; we formalize this classification in Section 2.1. It turns out that this is crucial to our study of bounded susceptibility, as follows:

For a fixed virulence β , the spreading time in the *SI* dynamics on star-like networks is small ($O(\log n)$, under fairly weak conditions on the non-core subgraph); in clique-like networks, the spreading time is $O(1)$ ¹. In case of the *SIS* epidemic, both clique-like networks and star-like networks experience exponential extinction times [8]. Thus, in a sense, *both star-like and clique-like networks are well suited for epidemic processes* – this behavior, however, *no longer holds under bounded susceptibility*.

1.1 An Overview of our Results

We refer to our modified epidemic models, incorporating bounded susceptibility, as the SI_{bs} and SIS_{bs} dynamics. As mentioned, we assume that in both models, a susceptible node j gets infected at a rate $\Phi(\beta|I_j|)$, where I_j are the infected neighbors of j ; in the SI_{bs} model, infected nodes remain so forever, while in the SIS_{bs} model, they recover to become susceptible at a rate 1. We assume that $\Phi(0) = 0$, and $\Phi(\cdot)$ is concave, non-decreasing, and bounded above by some $\Phi_{\max} = \Theta(1)$.

SI_{bs} Dynamics: In the SI_{bs} model, we are interested in finding the (expected) spreading time, under the assumption that both β and Φ_{\max} are $\Theta(1)$. In this context, we show:

- For clique-like networks, the spreading time switches from $O(1)$ under *SI* dynamics to $\Omega(\log n)$ under the SI_{bs} dynamics. In particular, for a $G(n, p)$ graph in the connected regime (i.e., $p \sim \omega(\log n/n)$), the spreading time switches from $o(1)$ to $\Omega(\log n)$ under the SI_{bs} dynamics.
- For star-like networks, the spreading time is $\Theta(\log n)$ under both *SI* and SI_{bs} dynamics (under some assumptions on the network structure). Thus, imposing bounded susceptibility does not significantly alter the epidemic behavior.

Thus, the *SI* model is highly sensitive to bounded susceptibility in clique-like networks – note that the spreading

time has switched from vanishing with network size to actually scaling with network size. On the other hand, it is insensitive in star-like networks.

A moment's thought suggests that this behavior is indeed reasonable – we have greatly reduced the (potential) susceptibility of high-degree nodes, which should hinder the spread of the epidemic. Furthermore, since high-degree nodes are more affected, it seems natural that bounded susceptibility should adversely affect dense graphs much more than sparse graphs. Surprisingly, the latter intuition is completely reversed in case of the SIS_{bs} model.

SIS_{bs} Dynamics: Recall that in the SIS_{bs} model, we want to find the (expected) extinction time under $\beta, \Phi_{\max} = \Theta(1)$ and recovery rate 1. In this context, we show the following:

- For clique-like networks, the extinction time is exponential (i.e., $\exp(\Omega(\text{poly}(n)))$) both with and without the bounded susceptibility assumption. In fact, this robustness property holds *for any network which contains a polynomial-sized subgraph which resembles a random $G(n, p)$ graph in the connected regime* (i.e., with $p \sim \omega(\log n/n)$).
- However for star-like networks, we now observe a sharp transition in extinction time. While star-like networks experience epidemics with exponential lifetimes under the *SIS* process [8, 9], we prove that this is not true in the case of SIS_{bs} infections. In particular, for any network which can be decomposed into a small *core*, plus an outer sparse subgraph, *the SIS_{bs} epidemic lifetime is sub-exponential for any virulence $\beta = \Theta(1)$.*

This reversal of results may seem paradoxical – there is however an intuitive explanation for this. Note that the absorbing states in the *SI* and *SIS* epidemics are exactly the opposite (all-infected vs. all-susceptible). Moreover, in the *SI* model, we want the epidemic to move rapidly towards the all-infected state, while in the *SIS*, we want to deter it from reaching the all-susceptible state. Rapid transition to the all-infected state in the *SI* dynamics is aided by having many high-degree nodes – this is hampered by imposing bounded susceptibility. On the other hand in the *SIS* dynamics, high degree nodes impede transition to the all-susceptible state – having many such nodes helps distribute this effect, but too few high-degree nodes means that these are now critical bottlenecks. Under bounded susceptibility, these bottleneck nodes are the most affected, resulting in a sharp transition in the epidemic behavior.

To make this intuition more transparent, in Section 2.2 we demonstrate the effect of the various dynamics in the clique and star networks, which are in a sense the extreme-case scenarios. Our main contribution in this work is to show how these phenomena extend to large classes of networks, occupying the space between these two extremes. In the process, we develop some structural lemmas relating the dynamics of epidemics to network properties, which may be of independent interest.

Technical Contributions. Most of our arguments involve standard Markov chain machinery, along with careful use of measure concentration. Some of the ideas we use are of fairly recent origin, in particular, the sharp expressions for extinction time in *SIS* models via embedding in an ergodic Markov chain (Lemma 1, which we adapt from Lemma 8 in [10]), and characterization of the *SI* spreading time in clique-like graphs, which is similar to results for first-passage percola-

¹Special cases of this result can be found in [2, 4]

tion in random graphs [4]. However, unlike [4], our results (Theorems 1, 3, 4 and 6) do not focus on a particular graph or generative model, but rather, relate the epidemic behavior to structural properties of the network. We believe these lemmas could be useful in studying other related settings. Finally, our results follow from elementary arguments, and hold for finite networks, with explicit error bounds in terms of network size.

We note that our classification of networks into clique-like and star-like does not cover all networks. In particular, there are networks which are a mix of the two, and where our results may not be tight; sharper techniques could perhaps be used to bridge this gap. Furthermore, our techniques do not immediately extend to other epidemic models, in particular, the SIR dynamics (see [12, 13] for details) – this is an interesting avenue for future exploration.

1.2 Bounded Susceptibility: Applications in Real-World Settings

Before we formally state our results, we first discuss three important real-life settings which exhibit bounded susceptibility, and point out applications of our results therein.

Social and Economic Networks. Perhaps the most compelling example of bounded susceptibility in network epidemics comes from online social networks. Sites such as Twitter and Facebook are built on a broadcast model – users upload content to the website, intending it to be broadcast to all neighbors. However, users are known to consume content from only a small set of neighbors – this is due to a combination of limited attention-spans and content filtering algorithms, and has been observed empirically in Twitter data [14, 15]. Epidemic models have been used to understand why content ‘goes viral’ on such media, however existing works do not account for bounded susceptibility.

Non-linear propagation models are also used in economics to model network externalities in the spread of opinions, technologies, etc. For example, the adoption of a new technology on a network is often subject to diminishing returns due to wearing out of novelty with increasing number of recommendations [16] – this naturally suggests a concave infection-profile, and further, bounded susceptibility.

Human-Disease Epidemiology. Network epidemic models are widely used to study human epidemics, and also computer viruses. Here again, bounded susceptibility is often a key feature of real epidemics, due to spread of awareness and adoption of preventive measures. Non-linear infection models have been used to capture the effect of risk perception in the spread of epidemics [17, 18], by combining the effect of infection spread and the protection measures taken by the susceptible individual. However these works focus on approximate and mean-field treatments for the problem.

Our work provides both intuition and technical tools for studying the effects of bounded susceptibility on epidemics. One important application of our findings is in the design of vaccination and quarantining policies on networks – such policies help reduce the susceptibility of nodes to infection, but usually nodes retain some susceptibility. A closely related design question is of distributing antidotes in a network [11] – there, the epidemic is controlled by increasing the recovery rate of a node, whereas our control is via limiting the susceptibility of a node.

Distributed Computing on Networks. A third application of epidemic processes is as a primitive for designing low-complexity distributed algorithms. The *SI* dynamics is the basis of the *flooding algorithm*, which is used for broadcast messaging in networks, distributed database synchronization, etc [19]. The *SIS* dynamics is sometimes used to ensure persistent storage in distributed storage/sensing applications with disk failures [20]. Existing work focuses on using the standard *SI/SIS* dynamics, or modifies them to impose *bounded influence*, i.e., a bound on the upload bandwidth of a node.

However many of the above settings, in reality, exhibit bounded susceptibility. Modern distributed/P2P systems are based on a virtual network architecture, where nodes upload content to a central location, from where it is downloaded by its neighbors. Thus, broadcasting is easy, but downloading is bandwidth constrained. This is also true for sensor networks – a node needs to broadcast a message once, but it is limited in how many neighbors it can listen to, due to a fixed download bandwidth. Thus, bounded susceptibility is natural in these settings. We discuss all these applications in Section 2.4, after presenting our results.

1.3 Related Work

Epidemic processes on networks have been studied across many disciplines; readers interested in more details regarding epidemics are referred to several excellent books on the subject [1, 12, 21, 22]. Specifically, there is a vast literature to characterize spreading time in various contexts for *SI* processes [1, 2, 3, 4], and phase transitions/extinction time for *SIS* processes [6, 7, 8, 9, 10, 11]. Phase transitions for *SIR* processes are available in [13].

Non-linear epidemic models have been explored in the past in various contexts. A large body of work is devoted to settings with *bounded influence*, where the total rate at which a node infects its neighbors is bounded [23, 24], but the infection-profile $\Phi(\cdot)$ is still linear. One popular model with non-linear $\Phi(\cdot)$ is the Bootstrap Percolation, and related variants [25, 26]. Several authors have also studied non-linear infection models at the population level using mean field approximations [17, 18] – these works show the existence of epidemic thresholds under different non-linear models. However, to the best of our knowledge, ours is the first work that explicitly characterizes epidemic behavior under bounded susceptibility.

2. MAIN RESULTS AND DISCUSSION

We now present and discuss our main results. Our general theorems are presented in Section 2.3, followed by a discussion of their application in Section 2.4; before that, in Section 2.2, we specialize our results for the clique and star networks, and give an outline of the proofs for these cases. The complete proofs are somewhat technical – we outline the main proof ideas after each theorem, deferring complete proofs to Section 3.

2.1 System Model

The SI/SIS Dynamics

The *SI*, *SIS* dynamics models we use are standard in literature [8, 10]. We consider a graph $G(V, E)$ with n nodes (vertices). The various processes (*SI*, *SI_{bs}*, *SIS*, *SIS_{bs}*) evolve on this graph in continuous time. Associated with each node

$i \in \{1, 2, \dots, n\}$ is a random process $X_i(t) \in \{0, 1\}$. Here, ‘0’ corresponds to the susceptible state and ‘1’ corresponds to the infected state (the entire state vector is denoted by $\mathbf{X}(t)$). For the *SI* and *SIS* models, nodes in state ‘1’ (infected nodes) attempt to infect a neighbor *independently* at a rate β , i.e., after a i.i.d random interval drawn from an *Exponential*(β) distribution. In the *SI* dynamics, an infected node never recovers. In contrast in the *SIS* dynamics, an infected node returns back to the susceptible state at rate 1 (i.e., after an *Exponential*(1) time).

We assume that the infection starts off at a single node, arbitrarily chosen. This is natural for the *SI* epidemic since we want to study spreading-time. For the *SIS* epidemic, our results generalize for any constant-sized initial infected set. For the *SI* dynamics, the only absorbing state is the *all-infected* state (i.e., $X_i = 1 \forall i$), and the spreading-time is defined as $\tau_s \triangleq \inf_t \{X_i(t) = 1 \forall i\}$. In contrast, for the *SIS* dynamics, the only absorbing state is the *all-susceptible* state (i.e., $X_i = 0 \forall i$), and the extinction-time is defined as $\tau_e \triangleq \inf_t \{X_i(t) = 0 \forall i\}$. We also assume throughout that β is a constant, independent of n ; however, our analysis extends for settings where β is a function of n .

Infection-Profiles and Bounded Susceptibility

A critical assumption in both *SI* and *SIS* models is that the infection rate for a susceptible node scales in proportion to the number of infected neighbors. In order to study the effect of non-linearity in the infection-rate, we introduce the notion of an infection-profile: we assume that a susceptible node j is infected at a rate $\Phi(\beta|I_j|)$, where I_j is the number of infected neighbors of node j (at some time t ; we suppress the dependence on t for ease of notation). Thus, in both the *SI*_{bs} and the *SIS*_{bs} dynamics, we have:

$$X_i : 0 \rightarrow 1 \quad \text{at rate } \Phi \left(\beta \sum_{(i,j) \in E} X_j \right),$$

and in the *SIS*_{bs}, we additionally have $X_i : 1 \rightarrow 0$ at rate 1.

$\Phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ represents the overall effect that the infected neighbors together have on a susceptible node. We assume that $\Phi(0) = 0$, and that $\Phi(\cdot)$ is non-decreasing, concave and bounded. Note that these assumptions imply that $\Phi_{\max} := \sup_{x \in \mathbb{R}} \Phi(x)$ exists, and $\Phi(\beta) > 0 \forall \beta > 0$.

Clique-Like and Star-Like Networks

We now formalize the two classes of networks, which we introduced in Section 1. Intuitively, clique-like networks are those which are dense (i.e., where $|E| = \omega(n)$) and have edges evenly distributed across the cuts of the network. One way to formalize this is via the expansion properties of the underlying graph. There are different ways to quantify the expansion of a graph. For example, the *edge expansion* of a graph is characterized by the *generalized isoperimetric constant*:

$$\eta_m^e(G) := \inf_{S \subseteq V, |S| \leq m} \frac{|\delta(S)|}{|S|}, \quad 0 < m \leq \lfloor n/2 \rfloor,$$

where, for any set $A \subseteq V$, $\delta(A) = \{(i, j) \in E \text{ s.t. } i \in A, j \notin A\}$ are the edges in the cut defined by A . The case of $m = n/2$ is referred to as the *isoperimetric/Cheeger constant*:

$$\eta(G) := \inf_{S \subseteq V, |S| \leq n/2} \frac{|\delta(S)|}{|S|}.$$

We can also define an alternate notion of isoperimetry based on *vertex expansion*, as follows:

$$\eta_m^v(G) := \inf_{S \subseteq V, |S| \leq m} \frac{|\Gamma(S)|}{|S|}, \quad 0 < m \leq \lfloor n/2 \rfloor,$$

where $\Gamma(S)$ is the *neighborhood* of S , i.e., the set of nodes in S^C with at least one neighbor in S . It is easy to check that for any fixed m , we have $\eta_m^e(G) \geq \eta_m^v(G)$.

The isoperimetric constants characterize bottleneck sets, i.e., those containing the smallest fraction of the potential edges in the cut. For a clique, it is easy to check that $\eta(G) = \Theta(n)$, which is the highest possible. To show that epidemic properties hold over a large class of graphs, we want to admit graphs with smaller isoperimetric constants. This motivates the following characterizations for clique-like networks:

DEFINITION 1 (CLIQUE-LIKE NETWORKS). *We define a graph G to be clique-like if it satisfies one of the following:*

- (A) *If $\eta(G) = \Omega(\log n)$.*
- (B) *If for some $m = n^\alpha$, $\alpha \in (0, 1)$, we have $\eta_m^v(G) = \omega(1)$.*

The two definitions arise from analyzing the two epidemic dynamics on graphs. Although not equivalent, they both characterize networks with good expansion properties. To demonstrate their wide applicability, we show that random (*Erdős-Rényi*) graphs in the connected regime (i.e., $G(n, p)$ with $p \sim \Omega(\log n/n)$) are clique-like under both definitions.

On the other hand, star-like networks intuitively are those which can be partitioned into a *dense core* and a *sparse periphery*. Here, the natural notions of density turn out to be the average degree $d_{\text{avg}}(G)$, and the spectral radius $\rho(G)$, i.e., the largest eigenvalue of the adjacency matrix, which is closely related to the node degrees (in particular, $d_{\text{avg}}(G) \leq \rho(G) \leq d_{\text{max}}(G)$). Formally, we write $V = V_c \cup V_p$, where V_c denotes the core and V_p the periphery. As for clique-like networks, we have two characterizations:

DEFINITION 2 (STAR-LIKE NETWORKS). *We define a graph G to be star-like if it satisfies one of the following:*

- (A) *If $|V_p| = \Omega(n)$ and the average degree of nodes in the periphery is $O(1)$. In addition, the diameter of the graph is $O(\log n)$.*
- (B) *If $|V_c| = O(\text{polylog } n)$ and the subgraph $G(V_p)$ induced by the periphery has spectral radius $\rho(G(V_p)) = o(1)$.*

Again, the class of graphs that are star-like by the two definitions above are not equivalent but broadly include networks which have a sparse periphery and a small (possibly dense) core. We derive our results for *SI* dynamics based on definition (A) of clique-like and star-like networks, while the results for *SIS* dynamics are based on definition (B) for both the networks.

2.2 Two Examples: The Clique and the Star

Before presenting our full results, we first illustrate them by considering our two representative graphs – the clique and the star. This comparison brings out the dichotomy between the two network classes. Further, it allows us to build intuition for our main results, since these special cases admit simple proofs. In all these results, we use standard notation (O , Θ , Ω , o and ω) to characterize the scaling behavior with the graph size, n . We also use the notation $[n] \triangleq \{1, 2, \dots, n\}$, and $\mathcal{H}_n = \sum_{i=1}^n \frac{1}{i}$.

SI – SI_{bs} Dynamics

Consider first the SI/SI_{bs} epidemic on K_n , the clique on n nodes. We now show that the spreading-time is *vanishingly small* under the SI dynamics, while it *scales with n* under the SI_{bs} dynamics.

PROPOSITION 1. *For the SI/SI_{bs} epidemic on K_n starting at an arbitrary node,*

1. *Under the SI dynamics, $\mathbb{E}[\tau_s] = \Theta(\frac{\log n}{n})$.*
2. *Under the SI_{bs} dynamics $\mathbb{E}[\tau_s] = \Theta(\log n)$.*

PROOF. The first part is known from previous results [4] and follows from standard Markov-chain arguments. By denoting $N(t)$ to be the number of infected nodes at time t , and analyzing the steady-state of the resulting (one-dimensional) CTMC, we immediately get the result.

Next, for the bounded-susceptibility spreading process, SI_{bs}, we again consider the process $\{N(t)\}$, i.e., the total number of infected nodes. $\{N(t)\}$ has transition rates as follows:

$$z \rightarrow z + 1 \quad \text{at rate } \Phi(\beta z)(n - z) \quad \forall 1 \leq z \leq n - 1,$$

and hence the expected spreading-time is given by:

$$\mathbb{E}[\tau_s] = \sum_{k=1}^{n-1} \frac{1}{\Phi(\beta k)(n - k)} \geq \sum_{k=1}^{n-1} \frac{1}{(n - k)\Phi_{max}} = \frac{\mathcal{H}_{n-1}}{\Phi_{max}},$$

where the inequality follows since $\Phi(\beta k) \leq \Phi_{max}$. From the same equation, we also have:

$$\mathbb{E}[\tau_s] = \sum_{k=1}^{n-1} \frac{1}{\Phi(\beta k)(n - k)} \leq \frac{\mathcal{H}_{n-1}}{\Phi(\beta)},$$

since Φ is a non-decreasing function. Finally, since both $\Phi(\beta)$ and Φ_{max} are $\Theta(1)$, and using $\ln(n+1) \leq \mathcal{H}_n \leq \ln(n) + 1$, we get the second assertion. \square

Thus there is a sharp change in the spreading-time between the SI and SI_{bs} epidemic processes in cliques. This is clearly visible in Figure 1, where we simulate and plot the spreading-time under the SI and SI_{bs} dynamics on a clique, as a function of network size n . Further, in Theorem 8 in the Appendix, we derive concentration bounds showing that the above result holds with high probability.

On the other hand, the same behavior is not seen in a star network, as we show via our next result:

PROPOSITION 2. *In a star network with a single hub and $n - 1$ leaves, under both the SI and SI_{bs} epidemics, we have:*

$$\mathbb{E}[\tau_s] = \Theta(\log n)$$

PROOF. Consider first the SI process on a star network. We denote the hub as node 1 and define:

$$\tau_1 = \inf_{t \geq 0} \{\text{Node 1 is infected}\},$$

i.e., the first time that the hub gets infected. Since the epidemic starts at a single (arbitrary) node, we can split the spreading-time as: $\tau_s = \tau_1 + (\tau_s - \tau_1)$. Next, we have that $\mathbb{E}[\tau_1] = O(1)$ and $\tau_s - \tau_1$ corresponds to the maximum of $m \in \{n - 2, n - 1\}$ i.i.d random variables, each drawn from a *Exponential*(β) distribution; thus $\mathbb{E}[\tau_s - \tau_1] = \frac{\mathcal{H}_m}{\beta} \sim \Theta(\log n)$. Finally, note that no susceptible node ever has more than 1 infected neighbor – hence the SI_{bs} model is identical to the SI model with $\beta' = \Phi(\beta) = \Theta(1)$. Therefore, the result also holds for SI_{bs} epidemics. \square




Figure 1: Simulated epidemic spreading-time in a clique under SI and SI_{bs}, averaged over 500 runs for $\beta = 0.1$ and $\Phi_{max} = 0.5$. Note that $\mathbb{E}[\tau_s]$ switches from $o(1)$ under the SI dynamics, to $\Omega(\log n)$ under SI_{bs} (decreasing to increasing on log – log plot).

SIS – SIS_{bs} Dynamics

Next, we consider the SIS/SIS_{bs} epidemic. Here, it is known [8] that the expected extinction-time is exponential in cliques – we will prove that it is exponential even for the SIS_{bs} dynamics.

We first state and prove a lemma that gives a closed-form expression for expected time to absorption in a birth-death Markov chain. The lemma is adapted from [10, Lemma 8] – it re-obtains the result for SIS epidemic on cliques, and also extends to our new result for the SIS_{bs} epidemic.

Consider a birth-death CTMC $\{U(t)\}$ on $\{0, 1, \dots, n\}$, with the following transition rates:

$$\begin{aligned} i \rightarrow i + 1 & \quad \text{at rate } q_{i,i+1} \quad \forall 1 \leq i \leq n - 1, \\ i \rightarrow i - 1 & \quad \text{at rate } q_{i,i-1} \quad \forall 1 \leq i \leq n. \end{aligned}$$

Note that 0 is the only absorbing state for the above process. Let $T(i, j) := \inf\{t : U(t) = j \mid U(0) = i\}$. Then we have:

LEMMA 1 (LEMMA 8 FROM [10]).

$$\mathbb{E}[T(1, 0)] = \frac{1}{q_{1,0}} \left(1 + \sum_{k=2}^n \prod_{i=1}^{k-1} \frac{q_{i,i+1}}{q_{i+1,i}} \right).$$

For the sake of completeness, we provide a proof of this result in the Appendix. We now use this to get lower bounds on the expected extinction-time in cliques.

PROPOSITION 3. *For the SIS/SIS_{bs} epidemic on clique K_n starting at an arbitrary node, under both the SIS and SIS_{bs} dynamics, $\log(\mathbb{E}[\tau_e]) = \Omega(n)$.*

PROOF. We prove the statement for the SIS_{bs} dynamics; the claim for the SIS dynamics follows similarly (and is also known from previous work [8]). As in Proposition 1, it is sufficient to consider the total number of infected nodes $N(t)$ at time t . $\{N(t)\}$ has transition rates as follows:

$$\begin{aligned} z \rightarrow z + 1 & \quad \text{at rate } \Phi(\beta z)(n - z) \quad \forall 1 \leq z \leq n - 1, \\ z \rightarrow z - 1 & \quad \text{at rate } z \quad \forall 1 \leq z \leq n. \end{aligned}$$

Let $k_0 = \inf\{k \in \mathbb{N} : \Phi(\beta k) \geq \Phi_{max}/2\}$. Since Φ is concave and $\Phi(0) = 0$, we have $\Phi(\beta k) \geq \frac{k}{k_0} \Phi(\beta k_0) \geq \frac{k}{k_0} \frac{\Phi_{max}}{2} \quad \forall k \leq$

k_0 . Note that $k_0 = \Theta(1)$. Now, from Lemma 1, we have:

$$\begin{aligned}
\mathbb{E}[\tau_e] &= 1 + \sum_{k=2}^n \prod_{i=1}^{k-1} \frac{\Phi(\beta i)(n-i)}{i+1} \\
&\geq 1 + \sum_{k=2}^{k_0} \prod_{i=1}^{k-1} \frac{\Phi_{max}}{2k_0} \frac{i(n-i)}{i+1} \\
&\quad + \sum_{k=k_0+1}^n \prod_{i=1}^{k_0-1} \frac{\Phi_{max}}{2k_0} \frac{i(n-i)}{i+1} \prod_{j=k_0}^{k-1} \frac{\Phi_{max}}{2} \frac{(n-j)}{j+1} \\
&= 1 + \sum_{k=2}^{k_0} \frac{(k-1)!}{n} \binom{n}{k} \left(\frac{\Phi_{max}}{2k_0} \right)^{k-1} \\
&\quad + \frac{k_0!}{n k_0^{k_0}} \sum_{k=k_0+1}^n \binom{n}{k} \left(\frac{\Phi_{max}}{2} \right)^{k-1} \\
&\geq \frac{2k_0!}{n \Phi_{max} k_0^{k_0}} \sum_{k=1}^n \binom{n}{k} \left(\frac{\Phi_{max}}{2} \right)^k \\
&= \frac{2k_0!}{n \Phi_{max} k_0^{k_0}} \left(\left(1 + \frac{\Phi_{max}}{2} \right)^n - 1 \right).
\end{aligned}$$

Since Φ_{max}, k_0 are $\Theta(1)$, $\log(\mathbb{E}[\tau_e]) = \Omega(n)$. \square

The more interesting phenomenon in the SIS model occurs in the star network. As in the clique, it is known [8, 9] that the expected extinction-time in star networks is exponential. However, under the SIS_{bs} dynamics, we show that it is sub-exponential. More precisely, we show that the epidemic survives in a star network for a time that is *at most polynomial* in n .

PROPOSITION 4. *For an SIS_{bs} epidemic originating at any node in a star graph, $\mathbb{E}[\tau_e] = O(n^{\Phi_{max}+1})$.*

This change from exponential to polynomial infection time is clearly seen in simulation results shown in Figure 2. The proof of this proposition is somewhat technical, so we provide only a proof sketch here. The high-level proof structure is similar to existing analysis of the SIS epidemic on a star [8, 9] – the novelty lies in showing that bounded susceptibility causes a sharp transition in the epidemic lifetime.

PROOF OUTLINE. We consider cycles of the epidemic evolution – each cyclical epoch starts off with the hub infected, then recovering to become susceptible, and finally getting reinfected, thereby starting the next cycle. The time from the start of an epoch to when the hub first becomes susceptible is an i.i.d *Exponential*(1) variable – at this time, some subset of leaf nodes are infected. Subsequently, either the hub is re-infected (i.e., a new epoch starts), or all these leaf nodes recover before the hub is re-infected (i.e., the epidemic becomes extinct). *The crux of the proof lies in estimating the probability that, under the SIS_{bs} model, the leaf nodes recover before the hub is reinfected.* Given this, we then show that the total number of such epochs is stochastically dominated by a geometric random variable, whose mean is polynomial in n – this gives the above result. \square

In Section 2.3, we state a much more general version (Theorem 6) of this result, extending it to star-like networks. The complete proof of both results are given in Section 3.

2.3 Main Results

We now state our main results, which generalize results from the previous section to clique-like and star-like graphs.




Figure 2: Simulated epidemic extinction-time in a star under SIS and SIS_{bs} , averaged over 500 simulation runs for $\beta = 0.1$ and $\Phi_{max} = 0.5$. Note that $\mathbb{E}[\tau_e]$ switches from exponential under SIS to polynomial under SIS_{bs} (polynomial to linear on log – log plot).

2.3.1 $SI - SI_{bs}$ Dynamics

Recall that, for a graph G , the Cheeger constant is given by $\eta(G) = \inf_{S \subset V, |S| \leq n/2} \frac{|\delta(S)|}{|S|}$. We now have the following theorem, which allows us to bound the spreading-time of the SI epidemic in clique-like graphs:

THEOREM 1. *Given underlying graph G with Cheeger constant $\eta(G)$, the spreading-time for an SI epidemic with virulence β originating at a single arbitrary node, obeys:*

$$\mathbb{E}[\tau_s] \leq \frac{(2 + o(1)) \log(n)}{\beta \eta(G)}.$$

The proof of this theorem, presented in Section 3, is similar to that of Proposition 1. Note that for the clique, we have $\eta(G) = n/2$. Plugging this into the above bound immediately gives back Proposition 1. However the theorem captures in finer detail the impact of expansion (as captured by $\eta(G)$) on spreading-time. As a corollary to the theorem, we have the following result for clique-like networks:

COROLLARY 1. *In a network that is clique-like as per Definition 1(A) (i.e., $\eta(G) = \Omega(\log n)$), the expected spreading time of an SI epidemic is $O(1)$.*

We note also that although Theorem 1 bounds the *expected* spreading-time, we can extend it to get a concentration bound for the spreading-time; refer to Theorem 7 in the Appendix for details. To see how the result can be used to characterize the spreading-time in a variety of clique-like graphs, we present the following corollary for the case of dense random graphs, i.e., Erdős-Rényi graphs in the connected regime:

COROLLARY 2. *Let G be an Erdős-Rényi graph $G(n, p)$, with $p > \frac{32 \log n}{n}$. Then with probability at least $1 - \frac{1}{n^2}$, we have $\eta(G) \geq \frac{np}{4}$. Therefore, for an SI epidemic,*

$$\mathbb{E}[\tau_s] \leq \frac{(8 + o(1)) \log n}{\beta np}$$

with probability at least $1 - \frac{1}{n^2}$.

Note that by the Borel-Cantelli Lemma, the above bound on the expected spreading-time holds almost surely as $n \rightarrow \infty$. The proof is provided in Section 3.

Putting the pieces together, we can now compare the SI and SI_{bs} processes in dense random graphs:

THEOREM 2. *Let G be an Erdős-Rényi graph $G(n, p)$, with $p > \frac{32 \log n}{n}$. For the SI/SI_{bs} epidemic on G , starting at an arbitrary node, we have:*

1. *Under the SI dynamics, $\mathbb{E}[\tau_s] = O\left(\frac{\log n}{np}\right)$ with probability at least $1 - \frac{1}{n^2}$.*
2. *Under the SI_{bs} dynamics $\mathbb{E}[\tau_s] = \Omega(\log n)$.*

PROOF. The first claim follows from Corollary 2 and Theorem 1. For the second, note that given any graph G , and G' obtained by adding edges to G , $\tau_s(G') \leq_{st} \tau_s(G)$ – this can be shown via a standard coupling argument. Thus, for any graph G , we can always lower bound the spreading-time under the SI_{bs} dynamics by that on the clique K_n . The second claim thus follows from Proposition 1. \square

Next we turn to the case of star-like graphs. To characterize SI epidemics on such networks, we first state a lemma that gives a lower bound on $\mathbb{E}[\tau_s]$ in terms of average degree:

LEMMA 2. *Consider a graph G with $|V| = n$, and an SI epidemic on G , starting at some node v_1 . Let $A \subseteq V$ be a subset of nodes containing v_1 , and let \hat{d}_{Ac} be the average degree of the nodes in $V \setminus A$, i.e., $\hat{d}_{Ac} = \frac{\sum_{v \notin A} d(v)}{n - |A|}$. Then:*

$$\mathbb{E}[\tau_s] \geq \frac{\log(n - |A|)}{\beta \hat{d}_{Ac}}$$

For example, for the star graph with v_1 chosen to be the hub, we have $\hat{d}_{\{v_1\}^c} = 1$, and so we get back the lower bound in Proposition 2. We now use the above lemma to show that imposing bounded susceptibility does not significantly change the spreading-time in star-like networks:

THEOREM 3. *Consider a network G of diameter $O(\log n)$ that is star-like by Definition 2(A). In other words, there exists a partition of the nodes into a core V_c and periphery V_p , parametrized by:*

- Core size: $|V_c| = m$ is at most a constant fraction of nodes (i.e. $m \leq cn$ for some $c < 1$). In other words, $|V_p| = \Omega(n)$.
- Sparse periphery: $\hat{d}_p = O(1)$, i.e., the average degree of nodes in the periphery is constant (or lesser).

Then both under the SI and SI_{bs} dynamics, $\mathbb{E}[\tau_s] = \Theta(\log n)$.

Note that the sparsity constraint is on the average degree of nodes in V_p , not in the subgraph induced by V_p , i.e., it also accounts for the edges in the cut $\delta(V_p)$.

PROOF. First consider the SI dynamics, starting at some arbitrary node v_1 . We want to derive a lower bound on the spreading-time using Lemma 2 – to do so, we choose $A = V_c \cup \{v_1\}$, and hence $A^c = V_p \setminus \{v_1\}$. Now using Lemma 2 and the fact that $|V_p| \sim \text{poly}(n)$, and that both β and \hat{d}_p are $\Theta(1)$ we have that $\mathbb{E}[\tau_s] = \Omega(\log n)$.

To obtain an upper bound on the spreading-time, we under-dominate the spread by considering epidemic spread on the shortest-path spanning tree of the network rooted at the initial infected node. The resulting depth is at most the

diameter of the tree. Thus, the spreading-time is stochastically upper bounded by the maximum of n i.i.d random variables distributed as sum of $\text{diam}(G)$ number of exponential random variables with rate β . Finally, using standard concentration results for the sum of exponentials (refer [3]), we have that $\mathbb{E}[\tau_s] = O(\log n)$.

For SI_{bs} epidemic, we can first over-dominate the epidemic spread by converting the network into a clique – we have from Proposition 1 that the time taken to infect all nodes is $\Omega(\log n)$. On the other hand, we can again under-dominate the spread by considering an SI_{bs} process on the minimum spanning tree which is equivalent to an SI process of rate $\Phi(\beta)$. Consequently, $\mathbb{E}[\tau_s] = \Theta(\log n)$. \square

2.3.2 $SIS - SIS_{bs}$ Dynamics

Next we look at the $SIS - SIS_{bs}$ dynamics on clique-like and star-like graphs. Recall that in this case, we are interested in the extinction-time τ_e . In particular, we are interested in characterizing networks in which the epidemic is *persistent*, i.e. the extinction-time is exponential in n (i.e., $\mathbb{E}[\tau_e] \sim \Omega(e^{\text{poly}(n)})$).

Existing works [8, 9, 10] demonstrate that the virulence thresholds for short-lived ($\mathbb{E}[\tau_e] \sim O(\log n)$) SIS epidemics depend on the spectral radius, $\rho(G)$, while thresholds for persistent epidemics are related to the generalized isoperimetric constants. Though the results for the SIS epidemic depend on edge-expansion isoperimetric constant $\eta_m^e(G)$, the vertex-expansion isoperimetric constant $\eta_m^v(G)$ turns out to be the appropriate notion for the SIS_{bs} epidemic, as follows:

THEOREM 4. *Consider an SIS_{bs} epidemic on graph G , with virulence β and infection-profile $\Phi(\cdot)$. Suppose for $m = n^\alpha$ for some $0 < \alpha < 1$, we have that:*

$$\Phi(\beta) \eta_m^v(G) > 1,$$

then $\mathbb{E}[\tau_e] = e^{\Omega(n^\alpha)}$, i.e., the SIS_{bs} epidemic is persistent.

PROOF OUTLINE. The proof is similar to that of Proposition 3. In particular, we construct an under-dominating CTMC for the number of infected nodes under SIS_{bs} . Next, we use Lemma 1 to get an expression for the extinction-time, and show that the bound is indeed exponential in n if $\Phi(\beta) \eta_m^v(G) > 1$. For complete proof, see Section 3. \square

An immediate consequence of this result is the following corollary for clique-like networks.

COROLLARY 3. *An SIS_{bs} epidemic is persistent in any network belonging to the class of clique-like networks as per Definition 1(B).*

Returning to the SIS epidemic, the strongest known condition for a persistent SIS epidemic is $\beta \eta_m^e(G) > 1$ [8]. Now since $\eta_m^e(G) > \eta_m^v(G)$, this implies that $\beta \eta_m^v(G) > 1$ is also sufficient for a persistent SIS epidemic, i.e., without bounded susceptibility. Theorem 4 shows that the condition for a persistent SIS_{bs} epidemic is closely related. Furthermore, as we show subsequently, the vertex-expansion condition is strong enough to recover all the results from [8], in particular for the $G(n, p)$ and power-law graphs.

We now use Theorem 4 to demonstrate that both SIS and SIS_{bs} epidemics have exponential extinction-times in dense random graphs:

THEOREM 5. *For both the SIS and SIS_{bs} epidemics on a $G(n, p)$ graph with $p > \frac{16 \log n}{n}$. For a large enough n , with probability at least $1 - \frac{1}{n^2}$ we have,*

$$\log(\mathbb{E}[\tau_e]) = \Omega(n).$$

PROOF OUTLINE. We show under the conditions of the theorem that $\Phi(\beta)\eta_m^v(G) > 1$ with probability $\geq 1 - \frac{1}{n^2}$ if $m = n^\alpha$ for any constant $\alpha \in (0, 1)$. The result then follows using Theorem 4. Complete proof is given in Section 3. \square

Note also that any graph that contains a large clique-like subgraph is prone to an exponentially long lasting infection, even under bounded susceptibility, if the initial infected node belongs to the subgraph. Clearly a clique, or a realization of an Erdős-Rényi graph in the connected regime experiences long lasting infection – however such dense random subgraphs appear in other graph families including ‘power-law’ or scale-free networks. We discuss this further in Section 2.4 (in particular, Corollary 5).

More generally, these results show that large dense networks experience persistent epidemics even under bounded susceptibility *due to the presence of large clique-like subgraphs*. This is in sharp contrast to existing work (for example, [9]) which suggests that persistent epidemics in such networks arise due to the presence of high-degree nodes (‘hubs’).

In fact, it turns out that the presence of such ‘hubs’ do not help – we took a step towards showing this in Proposition 4, where we showed that a star graph does not experience persistent SIS_{bs} epidemics. We now show that this is true for a much larger class of star-like graphs. As before, the node-set V is decomposed into core V_c and periphery V_p . We characterize the sparsity of the periphery via the *spectral radius* $\rho(G(V_p))$ of the periphery subgraph – recall that $d_{avg}(G(V_p)) \leq \rho(G(V_p)) \leq d_{max}(G(V_p))$. Now we have:

THEOREM 6. *Consider a network G , with the nodes partitioned into a core V_c and periphery V_p characterized by:*

- Core size: $|V_c| = m$.
- Periphery sparsity: *Parametrized by $\rho(G(V_p))$.*

Suppose $G(V_p)$ satisfies $\Phi(\beta)\rho(G(V_p)) = 1 - \epsilon$, for some $\epsilon > 0$. Then for an SIS_{bs} epidemic on G ,:

$$\log(\mathbb{E}[\tau_e]) = O(m \log n).$$

PROOF OUTLINE. The high-level proof structure is similar to the outline we gave for Proposition 2. We again consider cyclic epochs. An epoch starts with some node in the core infected, transitioning to all core nodes becoming susceptible, and finally ends when some node in the core gets reinfected beginning the next epoch. Next, we bound under the SIS_{bs} model the probability that, in any epoch, the epidemic dies out in all nodes of the periphery before any node in the core is reinfected. Finally, via stochastic domination by a geometric random variable, we get the above result. The complete proof is given in Section 3. \square

Though we state the above theorem for a general core size, it is most significant in settings the core is small. For example, if the core is of constant size, then $\mathbb{E}[\tau_e]$ *scales only polynomially with n* – note that this includes the star network (where $\rho(G(V_p)) = 0$, as the periphery is fully disconnected if we remove the hub) that we characterized in Proposition 4. More generally, the following result for star-like networks holds as a result of Theorem 6.

COROLLARY 4. *The extinction time of an SIS_{bs} epidemic is sub-exponential in any network that is star-like as per Definition 2(B).*

We note that we can also derive concentration bounds, showing that the bound on the order of the extinction-time holds with high probability (see Theorem 9 in the Appendix).

2.4 Discussion

We now show how our results apply to the settings we discuss in Section 1.2.

Social and Economic Networks. Epidemic models have been widely used to model the viral spread of content in social networks. In particular, the SIS model can be used to study the spread of ‘memes’, i.e., related pieces of content that re-appear cyclically. Furthermore, the exponential extinction-time of an epidemic is in a sense connected to long-lasting nature of popular memes.

However, the existence of popular memes is often regarded to be as a result of highly influential nodes, i.e., those with high degree. Such high-degree nodes are a distinctive feature of power-law graphs, which are often used to model such networks. Our results however suggest that persistent epidemics arise due to large random subgraphs rather than high-degree nodes. We now formalize this notion by applying Theorem 5 to power-law graphs.

We consider the case of power law graphs with exponent $2 < \gamma < 3$. For random power law graphs generated under the Chung-Lu model [27] with exponent $2 < \gamma < 3$, it is known that they contain a n^{ϵ_1} connected random graph [8], and further, that they contain an n^{ϵ_2} clique [28] for some $\epsilon_1, \epsilon_2 > 0$. The SIS epidemic is known to have exponential lifetime in such graphs [8]; using Theorem 5, we recover this behavior under bounded susceptibility:

COROLLARY 5. *For networks generated via the Chung-Lu model [27], if the power-law exponent $\gamma \in (2, 3)$ and $\Phi_{max} > 1 + v$ for some $v > 0$, the SIS_{bs} epidemic has exponentially long lifetime.*

Thus scale-free networks do experience persistent epidemics even under bounded susceptibility, due to the presence of large clique-like subgraphs. *This questions the belief that highly-connected nodes are most responsible for viral content.*

Epidemiology. Epidemic models help understand the dynamics of infectious diseases amongst human/animal populations, and the spread of computer viruses and worms. These models also help in the design of defense mechanisms against such epidemics – for example, policies for spreading *antidotes*, which increase the recovery rate in the SIS epidemic [11]. In this regard, our work helps in developing a theory for *partial quarantine policies*, wherein adopting preventative measures helps reduce (and bound) the susceptibility of a node. We make a distinction here between partial and *perfect* quarantine policies, wherein a node can be completely disconnected from the network. Achieving a perfect quarantine is often unfeasible and there is always some residual susceptibility.

In this regard, our results paint a somewhat bleak picture – *enforcing quarantine policies in dense networks has an insignificant effect on the epidemic lifetime*. More specifically, as long as there is some $\Theta(1)$ residual susceptibility, the epidemic lifetime remains exponential in clique-like networks.

Note also that a similar effect is seen in [11] (under linear susceptibility), where an effective policy needs to give an amount of antidote to a node which is proportional to its degree – this is clearly more difficult in dense graphs.

The story is somewhat more positive in star-like networks, where Theorem 6 shows that in a network with a small core and a sparse outer network, partial quarantining affects a large change in the extinction-time, reducing it to $\text{poly}(n)$. Star-like networks are relevant for studying the spread of an infectious disease through human transportation networks, which comprise a core of long distance air-links, overlaid on a slower road network. In case of an epidemic, a natural policy is to focus on safeguarding the core cities (via quarantining). For this, we have the following corollary:

COROLLARY 6. *Consider a network which can be decomposed into:*

- *A dense core of size $O(1)$, which can be fully connected.*
- *A sparse periphery, with $\rho(G(V_p)) = O(1)$ (no restriction on edges between periphery and core).*

Now suppose only core nodes have bounded susceptibility (via quarantining), and the unprotected periphery satisfies $\beta\rho(G(V_p)) \leq 1 - \epsilon$ for some $\epsilon > 0$. Then the epidemic is not persistent, i.e., $\mathbb{E}[\tau_e] = O(\text{poly}(n))$.

Note that given $\Omega(n)$ edges between the hub and the periphery, the SIS epidemic has an exponential lifetime. Thus, safeguarding hub cities does help reduce the epidemic lifetime dramatically. However, *distributing antidotes may be a much more effective policy than quarantining as it can effect much faster (upto $O(\log n)$) extinction-times.*

Distributed Computing Applications. Another important application of SI epidemics is as a primitive for algorithms like *flooding* [19], i.e., forwarding a message to all network nodes. In order to minimize the flooding time, one approach would be to make the underlying network clique-like – however our results indicate that under bounded susceptibility, the flooding time is $\Omega(\log n)$, which is identical to that in a star-like network (for example, a hierarchical network, with a constant-sized core and all other nodes connected to the core). On the other hand, our results also show that star-like networks are much easier to protect from viruses which spread like an SIS process. Considering the tradeoff between performance and robustness to infection, *star-like architectures appear more appropriate than clique-like architectures for distributed computing applications.*

In case of networks using SIS epidemics to ensure persistent data storage [20], the story is strikingly different from the above. Here, the presence of bounded susceptibility results in a sharp difference between epidemic lifetime in star-like and clique-like networks. However, our results also indicate that it is not necessary to have very high edge densities, or very well-designed networks in order to achieve long lifetimes – on the contrary, *as long as there exists a large-enough connected random subgraph, the epidemic lifetime is exponential even under bounded susceptibility.*

3. PROOFS

We now give detailed proofs for all our results. We use the standard stochastic dominance notation $\{\leq_{st}, \geq_{st}\}$ for relations between random variables. Specifically, $(X \leq_{st} Y) \implies \forall z, \mathbb{P}[X > z] \leq \mathbb{P}[Y > z]$.

Proofs from Section 2.3.1.

PROOF OF THEOREM 1. Let $N(t)$ be the number of infected nodes at time t and T_k be the first time at least k nodes are infected, i.e., $T_k = \inf\{t : N(t) = k\}$. Also, $\forall k \in [n]$, let \mathcal{V}_k denote the set of all subsets of V of size k , and let $S_k \in \mathcal{V}_k$ be the set of infected nodes at T_k .

Since $T_1 = 0$ and $T_n = \tau_s$, we can telescope and write $\tau_s = \sum_{k=1}^{n-1} (T_{k+1} - T_k)$. Recall that in case of the clique in Proposition 1, we had that $\mathbb{E}[T_{k+1} - T_k] = \frac{1}{\beta k(n-k)}$. This is no longer true, but we can upper bound the time in terms of the isoperimetric constant $\eta(G)$, as follows.

$$\begin{aligned} \mathbb{E}[T_{k+1} - T_k] &= \sum_{s_k \in \mathcal{V}_k} \mathbb{E}[T_{k+1} - T_k | S_k = s_k] \mathbb{P}[S_k = s_k] \\ &= \mathbb{E} \left[\frac{1}{\beta |\delta(S_k)|} \right]. \end{aligned}$$

By the definition of $\eta(G)$,

$$|\delta(S_k)| \geq \begin{cases} \eta(G)k & \text{if } k < n/2 \\ \eta(G)(n-k) & \text{if } k \geq n/2. \end{cases}$$

Therefore, we have

$$\begin{aligned} \mathbb{E}[\tau_s] &\leq \sum_{k=1}^{\lfloor n/2 \rfloor} \frac{1}{\beta \eta(G)k} + \sum_{\lfloor n/2 \rfloor + 1}^{n-1} \frac{1}{\beta \eta(G)(n-k)} \\ &\leq 2 \sum_{k=1}^{\lfloor n/2 \rfloor} \frac{1}{\beta \eta(G)k} \leq \frac{(2 + o(1)) \log(n/2)}{\beta \eta(G)}. \end{aligned}$$

This completes the proof. \square

Next we show that a $G(n, p)$ in the connected regime has a large isoperimetric constant with high probability.

PROOF OF COROLLARY 2. Let graph G be a realization of a $G(n, p)$ random graph with $p > \frac{32 \log n}{n}$. We claim that

$$\mathbb{P}[\eta(G) < np/4] \leq \frac{1}{n^2}.$$

In other words, for all but a vanishing fraction of graphs (of measure $\leq 1/n^2$), we have the required bound on $\eta(G)$. To show this, we write

$$\begin{aligned} \mathbb{P}[\eta(G) < np/4] &= \mathbb{P} \left[\bigcup_{k \leq \lfloor n/2 \rfloor} \bigcup_{s_k \in \mathcal{V}_k} |\delta(s_k)| < nkp/4 \right] \\ &\leq \mathbb{P} \left[\bigcup_{k \leq \lfloor n/2 \rfloor} \bigcup_{s_k \in \mathcal{V}_k} |\delta(s_k)| < k(n-k)p/2 \right] \\ &\leq \sum_{k \leq \lfloor n/2 \rfloor} \sum_{s_k \in \mathcal{V}_k} \mathbb{P}[|\delta(s_k)| < k(n-k)p/2]. \end{aligned}$$

We can now use the following Chernoff bound – for any r.v. B taking values in $[0, 1]$, we have

$$\mathbb{P}[B < (1 - \delta)\mathbb{E}[B]] \leq \exp \left(\frac{-\mathbb{E}[B]\delta^2}{2} \right).$$

Since $|\delta(s_k)| \sim \text{Binom}(k(n-k), p)$,

$$\begin{aligned} \mathbb{P}[\eta(G) < np/4] &\leq \sum_{k=1}^{\lfloor n/2 \rfloor} \binom{n}{k} \exp \left(\frac{-1}{8} k(n-k)p \right) \\ &\leq n \binom{n}{1} \exp \left(\frac{-1}{8} (n-1)p \right). \end{aligned}$$

Using the condition $p > \frac{32 \log n}{n}$, we get our result. \square

Next, we turn to Lemma 2, wherein we derive a lower bound on the spreading-time in terms of the average degree of a graph. First, we need the following lemma:

LEMMA 3. Let $\{X_i\}_{i=1}^m$ be independent exponential random variables with rate a_i (i.e., $\text{Exponential}(a_i)$) such that $\sum_{i=1}^m a_i = A$, and let $\{Y_i\}_{i=1}^m$ be i.i.d exponential random variables with rate A/m . If $X_{\max} = \max\{X_1, X_2, \dots, X_m\}$ and $Y_{\max} = \max\{Y_1, Y_2, \dots, Y_m\}$, then $X_{\max} \geq_{st} Y_{\max}$.

PROOF. For stochastic dominance, we want to show that $\mathbb{P}[X_{\max} \leq x] \leq \mathbb{P}[Y_{\max} \leq x] \forall x > 0$. For any $x > 0$, define $Q_x(\mathbf{a}) = \prod_{i=1}^m (1 - \exp(-a_i x))$. We now show that $\mathbf{a}^* = (A/m, A/m, \dots, A/m)$ maximizes Q_x over the set $\mathcal{A} = \{\mathbf{a} : \sum_{i=1}^m a_i = A\}$. For $m = 2$, one can check that $\mathbf{a}^* = (A/2, A/2)$ is the unique maximum. Thus, $Q_x(\mathbf{a}^*) > Q_x(\mathbf{a})$ for any $\mathbf{a} \neq \mathbf{a}^*$. If $m > 2$, suppose that $\tilde{\mathbf{a}} \in \mathcal{A}$ and $\tilde{\mathbf{a}} \neq \mathbf{a}^*$, maximizes Q_x over \mathcal{A} . Then $\exists i, j$ such that $\tilde{a}_i \neq \tilde{a}_j$. Then, for $\hat{\mathbf{a}}$ equal to $\tilde{\mathbf{a}}$ except with the i, j elements replaced with $(\tilde{a}_i + \tilde{a}_j)/2$, we have $Q_x(\hat{\mathbf{a}}) > Q_x(\tilde{\mathbf{a}})$, which is a contradiction. Therefore, $\mathbf{a}^* = (A/m, A/m, \dots, A/m)$ maximizes Q_x over the set \mathcal{A} , and $\mathbb{P}[Y_{\max} \leq x] = Q_x(\mathbf{a}^*) \geq Q_x(\mathbf{a}) = \mathbb{P}[X_{\max} \leq x]$ for any $x > 0$. \square

Using this, we now prove Lemma 2.

PROOF OF LEMMA 2. Let the nodes be numbered from 1 to n such that $A = \{1, 2, \dots, |A|\}$ and let $\mathbf{d} = (d_1, d_2, \dots, d_n)$ be the degree sequence of G . At any time, the rate at which a susceptible node i is infected is upper bounded by βd_i . Therefore, $\tau_s \geq_{st} \max\{X_{|A|+1}, \dots, X_n\}$ where $\{X_i\}$ are independent exponential random variables with mean $1/\beta d_i$. By Lemma 3, $\tau_s \geq_{st} Y$, where Y is the maximum of $n - |A|$ i.i.d. exponential random variables with mean $\frac{n - |A|}{\sum_{i=|A|+1}^n \beta d_i} = \frac{1}{\beta \hat{d}_{Ac}}$.

$$\mathbb{E}[\tau_s] \geq \mathbb{E}[Y] = \frac{\mathcal{H}_{n-|A|}}{\beta \hat{d}_{Ac}} > \frac{\log(n - |A|)}{\beta \hat{d}_{Ac}}.$$

This completes the proof. \square

Proofs from Section 2.3.2.

PROOF OF THEOREM 4. We analyze a CTMC $Z(t)$ that is stochastically dominated by the total number of infected nodes $N(t)$ in the SIS_{bs} epidemic. This being true, we have that the expected time to absorption for $Z(t)$ is less than that of $N(t)$. We note that this proof structure is similar to that of Theorem 4.1 from [8]; however it is much simplified via the ergodic embedding technique of [10].

Note that for any set of nodes $A \subseteq V$ of size $|A| \leq m$ (for some $m = n^\alpha$), we have that the neighborhood of A has size at least $\eta_m^v(G)|A|$. Now suppose the set of infected nodes I satisfies $|I| \leq m$. Under bounded susceptibility, each node in $\Gamma(I)$ is infected at a rate which is $\geq \Phi(\beta)$. Using this observation, we define $Z(t)$ to have transition rates given by

$$\begin{aligned} z &\rightarrow z + 1 && \text{at rate } \eta_m^v(G)\Phi(\beta)z \quad \forall 1 \leq z \leq m - 1, \\ z &\rightarrow z - 1 && \text{at rate } z \quad \forall 1 \leq z \leq m, \end{aligned}$$

and all other transitions equal to 0. Note that $Z(t)$ is a finite state space CTMC on state space $\{0, 1, \dots, m\}$ – standard

coupling arguments (for example, see [12]) show that it is dominated by $N(t)$. Now using Lemma 1, we have

$$\mathbb{E}[\tau_e] \geq \mathbb{E}[\inf\{t : Z(t) = 0\}] = \sum_{k=0}^{m-1} (\eta_m^v(G)\Phi(\beta))^k.$$

So, if $\Phi(\beta)\eta_m^v(G) > 1$, we have $\log(\mathbb{E}[\tau_e]) = \Omega(n^\alpha)$. \square

Now we show that for the Erdős-Rényi graph in the connected regime, the condition for persistent SIS_{bs} epidemics is satisfied with high probability. Note that this proof also immediately gives persistence for the SIS epidemic.

PROOF OF THEOREM 5. Let G be a realization of a $G(n, p)$ graph. Fix $\alpha \in (0, 1)$, $m = n^\alpha$, and choosing some $\gamma > 1$, we define $b = \gamma/\Phi(\beta)$. We now prove that $\mathbb{P}[\eta_m^v(G) < b] \leq 1/n^2$ for $p = \frac{16 \log n}{n}$. Then by stochastic dominance, the statement is true for any $p > \frac{16 \log n}{n}$. First we have

$$\mathbb{P}[\eta_m^v(G) < b] \leq \sum_{i=1}^m \sum_{S: |S|=i} \mathbb{P}[|\Gamma(S)| < bi].$$

Defining $B(n, p, i) \sim \text{Binomial}(n - i, 1 - (1 - p)^i)$, we can write

$$\mathbb{P}[\eta_m^v(G) < b] \leq \sum_{i=1}^m \binom{n}{i} \mathbb{P}[B(n, p, i) < bi].$$

Note that $(1 - p)^i \leq 1 - ip + i^2 p^2 / 2 \forall i$. This implies that

$$\mathbb{E}[B] \geq (n - i)ip(1 - ip/2).$$

Defining $\delta = \left(1 - \frac{bi}{(n-i)ip(1-ip/2)}\right)$ and using Chernoff bound,

$$\begin{aligned} \mathbb{P}[B(n, p, i) < bi] &\leq \mathbb{P}[B(n, p, i) < (1 - \delta)\mathbb{E}[B(n, p, i)]] \\ &\leq \exp\left(\frac{-\mathbb{E}[B(n, p, i)]\delta^2}{2}\right) \\ &\leq \exp\left(\frac{-(n - i)ip(1 - ip/2)\delta^2}{2}\right). \end{aligned}$$

Note that $\frac{(n-i)}{n}(1 - ip/2)\delta^2 \rightarrow 1$ as $n \rightarrow \infty$. So for n large enough, $(n - i)(1 - ip/2)\delta^2 > n/2$ which implies

$$\begin{aligned} \mathbb{P}[\eta_m^v(G) < b] &\leq \sum_{i=1}^m \binom{n}{i} \exp\left(-\frac{1}{2}(n - i)ip(1 - ip/2)\delta^2\right) \\ &\leq \sum_{i=1}^m \frac{n^i}{i!} \exp(-i(np/4)) \\ &= \sum_{i=1}^m \frac{1}{i!} \exp(-i(np/4 - \log n)) < \frac{1}{n^2}. \end{aligned}$$

\square

Next we turn to Theorem 6. The high-level proof structure is similar to existing analysis for the SIS epidemics on stars [8, 9]. However, to study the effect of bounded susceptibility, we need to carefully characterize various inter-epoch events, which require controlling the size of the core and the sparsity of the periphery.

We need two crucial lemmas for the proof. The first lemma is a crude but general upper bound for SIS_{bs} extinction-time in *all graphs*,

LEMMA 4. For any network of size n , and for any arbitrary initial set of infected nodes, the expected extinction-time for SIS_{bs} dynamics has an upper bound given by

$$\mathbb{E}[\tau_e] \leq \frac{(1 + \Phi_{max})^n}{n\Phi_{max}}.$$

PROOF. We use the expression for $\mathbb{E}[\tau_e]$ from Lemma 1. For any susceptible node with an infected neighbor, we have that the infection rate is bounded by Φ_{max} . Further, at any time when the number of infected nodes $|I| \geq 1$, we can stochastically upper bound the infection by assuming that all susceptible nodes have an infected neighbor. Thus

$$\begin{aligned} \mathbb{E}[\tau_e] &\leq 1 + \sum_{k=2}^n \prod_{i=1}^{k-1} \frac{\Phi_{max}(n-i)}{i+1} \\ &= \frac{1}{n\Phi_{max}} \sum_{k=1}^n \binom{n}{k} \Phi_{max}^k \leq \frac{(1 + \Phi_{max})^n}{n\Phi_{max}}. \end{aligned}$$

□

We use Lemma 4 to control extinction-time in dense subgraphs. In case of sparse subgraphs, we can use a result from [8], which gives conditions for fast extinction of the SIS epidemic (and hence SIS_{bs} epidemic) in terms of $\rho(G)$.

LEMMA 5. [8, Theorem 3.1] Let $\rho(A)$ be the spectral radius of the adjacency matrix A . If $\beta < 1/\rho(A)$, then the probability that the epidemic has not died out by time t , given the initial condition $X(0)$, admits the following upper bound.

$$\mathbb{P}[X(t) \neq 0] \leq \sqrt{n\|X(0)\|_1} e^{-(1-\beta\rho(A))t},$$

where $\|X(0)\|_1 = \sum_{i=1}^n X_i(0)$.

With these two ingredients, we can now prove Theorem 6.

PROOF OF THEOREM 6. We define cyclical epochs, starting with some (arbitrary) set of nodes in the core V_c infected, transitioning to all nodes in V_c becoming susceptible, and then finally at least one node in V_c getting reinfected. This technique is similar those used in [8, 9] – however instead of a star, we now have to deal with *core and periphery subgraphs* and also *bounded susceptibility*.

Consider epoch j (for any $j \in \mathbb{N}$) starting at time $T_{0,j}$. We define $T_{1,j} := \inf\{t > T_{0,j} : \sum_{i \in V_c} X_i(t) = 0\}$, i.e., all nodes in V_c become uninfected. From Lemma 4, we have

$$\mathbb{E}[T_{1,j} - T_{0,j}] \leq \frac{(1 + \Phi_{max})^m}{m\Phi_{max}}, \quad (1)$$

where $m = |V_c|$. Note that, under bounded susceptibility, this *only depends on m and not on core/periphery topologies*.

Next, we define $T_{0,j+1}$ to be the first time after $T_{1,j}$ when either some node in V_c gets re-infected or all the nodes in V_p recover – in case of the former, the next epoch starts, while in case of the latter, the infection dies out (i.e., $T_{0,j+1} = \tau_e$).

It is clear that the random variable $T_{0,j+1} - T_{1,j}$ is stochastically dominated by ζ_1 , the time taken by all the nodes in V_p to completely recover if the periphery was isolated from the core at $T_{1,j}$. Now given that the infection profile Φ is concave, and $\Phi(0) = 0$, we have $\Phi(\beta x) \leq \Phi(\beta)x \quad \forall x \geq 1$ and so the number of infected nodes under the SIS_{bs} dynamics is stochastically dominated by that in an SIS infection process with rate $\Phi(\beta)$. Now given that $\Phi(\beta)\rho(G(V_p)) = 1 - \epsilon$ for some $\epsilon > 0$, where $\rho(G(V_p))$ is the spectral radius of the

periphery subgraph, we have from Lemma 5 that $\mathbb{P}(\zeta_1 > t) \leq \min\left\{ne^{-(1-\Phi(\beta)\rho(G(V_p)))t}, 1\right\}$. Thus

$$\begin{aligned} \mathbb{E}[T_{0,j+1} - T_{1,j}] &\leq \mathbb{E}[\zeta_1] = \int_0^\infty \mathbb{P}(\zeta_1 > t) dt \\ &\leq \int_0^\infty ne^{-(1-\Phi(\beta)\rho(G(V_p)))t} dt \\ &= \frac{n}{1 - \Phi(\beta)\rho(G(V_p))}. \end{aligned} \quad (2)$$

We now derive a lower bound for the probability q that the infection dies in the periphery $G(V_p)$ before any node in the core gets reinfected. Since each node in the core V_c is infected at a rate at most Φ_{max} , the time after $T_{1,j}$ for the core to get infected can be bounded from below by ζ_2 , an $Exponential(m\Phi_{max})$ random variable, where $m = |V_c|$. Defining $t_0 = \frac{\log n}{1 - \Phi(\beta)\rho(G(V_p))}$, we have

$$\begin{aligned} \mathbb{P}[\zeta_1 > \zeta_2] &= m\Phi_{max} \int_0^\infty \mathbb{P}(\zeta_1 > t) e^{-m\Phi_{max}t} dt \\ &\leq m\Phi_{max} \int_0^{t_0} e^{-m\Phi_{max}t} dt \dots \\ &\quad + m\Phi_{max} \int_{t_0}^\infty ne^{-(1-\Phi(\beta)\rho(G(V_p)) + m\Phi_{max})t} dt \\ &= 1 - e^{-m\Phi_{max}t_0} \dots \\ &\quad + \frac{nm\Phi_{max}e^{-(1-\Phi(\beta)\rho(G(V_p)) + m\Phi_{max})t_0}}{1 - \Phi(\beta)\rho(G(V_p)) + m\Phi_{max}} \\ &= 1 - \frac{1 - \Phi(\beta)\rho(G(V_p))}{1 - \Phi(\beta)\rho(G(V_p)) + m\Phi_{max}} n^{\frac{-m\Phi_{max}}{1 - \Phi(\beta)\rho(G(V_p))}} \\ &= 1 - \frac{\epsilon}{\epsilon + m\Phi_{max}} n^{\frac{-m\Phi_{max}}{\epsilon}}. \end{aligned}$$

The probability q that the infection dies before any node in the core gets infected can thus be bounded from below.

$$q \geq \mathbb{P}[\zeta_1 \leq \zeta_2] \geq \left(Cn^{\frac{m\Phi_{max}}{\epsilon}}\right)^{-1}, \quad (3)$$

where $C = 1 + m\Phi_{max}/\epsilon > 1$.

Finally, the extinction-time τ_e can now be bounded as

$$\tau_e \leq st \sum_{i=1}^{V+1} U_i,$$

where V is a geometric random variable with mean $\frac{1}{q}$ and U_i are i.i.d. random variables independent of V with mean equal to the sum of upper bounds (R.H.S) in Equations 1 and 2. Thus, using Equations 3, 1, 2, we have

$$\begin{aligned} \mathbb{E}[\tau_e] &\leq \frac{1}{q} \mathbb{E}[U_1] \\ &\leq Cn^{\frac{m\Phi_{max}}{\epsilon}} \left(\frac{(1 + \Phi_{max})^m}{m\Phi_{max}} + \frac{n}{1 - \Phi(\beta)\rho(G(V_p))} \right). \end{aligned}$$

Simplifying we get that $\log \mathbb{E}[\tau_e] = O(m \log n)$. □

4. CONCLUSION

We have studied the effect of bounded susceptibility on the SI and SIS epidemics. In particular, we showed that imposing this constraint brings out a sharp dichotomy between star-like and clique-like networks. Our results are applicable

in several real-life settings which exhibit bounded susceptibility – social networks, human-disease networks, infrastructure networks, etc. Along the way, we also develop a suite of lemmas relating graph structure and epidemic behavior – some extend/simplify existing results while others are novel. We hope these will prove useful in other settings.

Acknowledgements

This work was supported by NSF grants IIS-1017525, CNS-1320175 and ARO grant W911NF-11-1-0265. We thank the reviewers for their suggestions which helped in greatly improving the clarity of this paper.

5. REFERENCES

- [1] H. Kesten, “First-passage percolation,” in *From classical to modern probability*. Springer, 2003.
- [2] I. Benjamini and Y. Peres, “Tree-indexed random walks on groups and first passage percolation,” *Probability Theory and Related Fields*, 1994.
- [3] S. Banerjee, A. Gopalan, A. K. Das, and S. Shakkottai, “Epidemic spreading with external agents,” *IEEE Trans. Information Theory*, 2014.
- [4] S. Bhamidi, R. Van der Hofstad, and G. Hooghiemstra, “First passage percolation on the erds-renyi random graph,” *Combinatorics, Probability & Computing*, 2011.
- [5] T. M. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*. Springer, 1999.
- [6] J. O. Kephart and S. R. White, “Directed-graph epidemiological models of computer viruses,” in *IEEE Symposium on Security and Privacy*, 1991.
- [7] R. Pastor-Satorras and A. Vespignani, “Epidemic Dynamics in Finite Size Scale-Free Networks,” *Phys. Rev. E*, 2002.
- [8] A. J. Ganesh, L. Massoulié, and D. F. Towsley, “The effect of network topology on the spread of epidemics,” in *IEEE INFOCOM*, 2005.
- [9] N. Berger, C. Borgs, J. T. Chayes, and A. Saberi, “On the spread of viruses on the internet,” in *ACM-SIAM SODA*, 2005.
- [10] S. Banerjee, A. Chatterjee, and S. Shakkottai, “Epidemic thresholds with external agents,” in *IEEE INFOCOM*, 2014.
- [11] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, “How to distribute antidote to control epidemics,” *Random Structures & Algorithms*, 2010.
- [12] M. Draief and L. Massoulié, *Epidemics and rumours in complex networks, volume 369 of London Mathematical Society Lecture Notes*. Cambridge University Press, Cambridge, 2010.
- [13] M. Draief, A. Ganesh, and L. Massoulié, “Thresholds for virus spread on networks,” *The Annals of Appl. Prob.*, vol. 18, pp. 359–378, 2008.
- [14] N. O. Hodas and K. Lerman, “How visibility and divided attention constrain social contagion,” in *IEEE Privacy, Security, Risk and Trust*, 2012.
- [15] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, “Competition among memes in a world with limited attention,” *Scientific Reports*, 2012.
- [16] J. Leskovec, L. A. Adamic, and B. A. Huberman, “The dynamics of viral marketing,” *ACM Transactions on the Web (TWEB)*, 2007.
- [17] S. Funk, E. Gilad, C. Watkins, and V. A. A. Jansen, “The spread of awareness and its impact on epidemic outbreaks,” *PNAS*, 2009.
- [18] Q. Wu, X. Fu, M. Small, and X.-J. Xu, “The impact of awareness on epidemic spreading in networks,” *Chaos: Interdisciplinary Journal of Nonlinear Science*, 2012.
- [19] P. T. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massoulié, “From epidemics to distributed computing,” *IEEE Computer*, vol. 37, 2004.
- [20] D. Chakrabarti, J. Leskovec, C. Faloutsos, S. Madden, C. Guestrin, and M. Faloutsos, “Information survival threshold in sensor and p2p networks,” in *IEEE INFOCOM*, 2007.
- [21] R. M. Anderson and R. M. May, *Infectious Diseases of Humans Dynamics and Control*. OUP, 1992.
- [22] R. Durrett, *Random graph dynamics*. Cambridge university press, 2007.
- [23] D. Shah, “Gossip algorithms,” *Foundations and Trends in Networking*, 2009.
- [24] S. Sanghavi, B. Hajek, and L. Massoulié, “Gossiping with multiple messages,” in *IEEE INFOCOM*, 2007.
- [25] S. Janson, T. Luczak, T. Turova, and T. Vallier, “Bootstrap percolation on the random graph $g_{\{n,p\}}$,” *The Annals of Applied Probability*, 2012.
- [26] M. Lelarge, “Efficient control of epidemics over random networks,” in *Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems*, 2009.
- [27] F. Chung, L. Lu, and V. Vu, “Eigenvalues of random power law graphs,” *Annals of Combinatorics*, 2003.
- [28] S. Janson, T. Luczak, and I. Norros, “Large cliques in a power-law random graph,” *Journal of Applied Probability*, 2010.
- [29] P. Bremaud, *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag New York Inc., 2001.
- [30] O. Jouini and Y. Dallery, “Moments of first passage times in general birth–death processes,” *Mathematical Methods of Operations Research*, 2008.

APPENDIX

In this section, we give concentration results which show that there is a sharp contrast in spread/extinction times between the bounded susceptibility and the traditional models.

First however, we give a proof of Lemma 1 from Section 2.2. This proof is very similar to that of Lemma 8 in [10] – we include it here mainly for the sake of completeness.

PROOF OF LEMMA 1. Following [10], we embed this process in another CTMC $\{\tilde{U}(t)\}$ with same transition rates *except for an additional transition from 0 to 1 at rate $q_{0,1}$* . It is easy to see that $\{\tilde{U}(t)\}$ is now an ergodic CTMC. Recall that we defined $T(i, j) := \inf\{t : U(t) = j \mid U(0) = i\}$. Similarly, define $\tilde{T}(i, j) := \inf\{t : \tilde{U}(t) = j \mid \tilde{U}(0) = i\}$. Note that $T(1, 0)$ and $\tilde{T}(1, 0)$ have the same distribution, and therefore, $\mathbb{E}[T(1, 0)] = \mathbb{E}[\tilde{T}(1, 0)]$. Now, $\{\tilde{U}(t)\}$ is a finite state, irreducible CTMC and therefore, positive recurrent. Let π be its stationary distribution. Solving for π_0 gives:

$$\frac{1}{\pi_0} = 1 + \sum_{k=1}^n \prod_{i=0}^{k-1} \frac{q_{i,i+1}}{q_{i+1,i}}.$$

From standard ergodic Markov Chain theory [29], we have that $\mathbb{E}[\tilde{T}(0, 0)] = \frac{1/q_{0,1}}{\pi_0}$, and further, from our construction, $\mathbb{E}[\tilde{T}(0, 0)] = 1/q_{0,1} + \mathbb{E}[\tilde{T}(1, 0)]$. Thus:

$$\mathbb{E}[\tilde{T}(1, 0)] = \frac{1}{q_{0,1}} \left(\frac{1}{\pi_0} - 1 \right) = \frac{1}{q_{1,0}} \left(1 + \sum_{k=2}^n \prod_{i=1}^{k-1} \frac{q_{i,i+1}}{q_{i+1,i}} \right).$$

This completes the proof. \square

SI: Upper Bound in Clique-like Networks

THEOREM 7. *Under the conditions of Theorem 1, the spreading-time $\tau_s \leq \frac{(4+o(1))(\log n)^2}{\beta\eta(G)}$ with probability at least $1 - \frac{1}{n}$.*

PROOF. Define $\Delta(k) = \frac{2 \log n}{\beta\eta(G) \min(k, n-k)}$. Then we have:

$$\begin{aligned} \mathbb{P}[T_{k+1} - T_k > \Delta(k)] &= \sum_{s_k \in \mathcal{V}_k} \mathbb{P}[T_{k+1} - T_k > \Delta(k) \mid S_k = s_k] \mathbb{P}[S_k = s_k] \\ &= \mathbb{E} \left[\exp \left(-\frac{2\beta|\delta(S_k)| \log n}{\beta\eta(G) \min(k, n-k)} \right) \right] \leq \exp(-2 \log n). \end{aligned}$$

By telescoping, we have $\tau_s = \sum_{k=1}^{n-1} T_{k+1} - T_k$. Further:

$$\sum_{k=1}^{n-1} \Delta(k) = \sum_{k=1}^{n-1} \frac{2 \log n}{\beta\eta(G) \min(k, n-k)} \leq \frac{4(\log n)^2}{\beta\eta(G)},$$

Putting all the above together, we get:

$$\begin{aligned} \mathbb{P} \left[\tau_s > \frac{4(\log n)^2}{\beta\eta(G)} \right] &\leq \mathbb{P} \left[\sum_{k=1}^{n-1} T_{k+1} - T_k > \sum_{k=1}^{n-1} \Delta(k) \right] \\ &\leq \sum_{k=1}^{n-1} \mathbb{P}[T_{k+1} - T_k > \Delta(k)] \leq \frac{1}{n}. \end{aligned}$$

This completes the proof. \square

Therefore, in clique-like networks with $\eta(G) = \omega((\log n)^2)$, the spread time under the SI model is vanishingly small with high probability. Again, as in Proposition 1, we show that this is not the case with the SI_{bs} epidemic.

SI_{bs} : Lower Bound

THEOREM 8. *For any graph under SI_{bs} dynamics, $\tau_s > \frac{\log n}{2\Phi_{max}} w.p.$ at least $1 - \frac{1}{\sqrt{n}}$.*

PROOF. Since any susceptible node is infected by rate at most Φ_{max} , $T_{k+1} - T_k \geq_{st} Z_k$, where Z_k has *Exponential* $((n-k)\Phi_{max})$ distribution. Using Chernoff bound,

$$\begin{aligned} \mathbb{P} \left[\sum_{k=1}^{n-1} Z_k \leq t \right] &\leq e^{\Phi_{max} t} \prod_{k=1}^{n-1} \mathbb{E}[e^{-\Phi_{max} Z_k}] \\ &= e^{\Phi_{max} t} \prod_{k=1}^{n-1} \frac{k}{k+1} = \frac{e^{\Phi_{max} t}}{n}. \end{aligned}$$

Therefore,

$$\mathbb{P} \left[\tau_s \leq \frac{\log n}{2\Phi_{max}} \right] \leq \mathbb{P} \left[\sum_{k=1}^{n-1} Z_k \leq \frac{\log n}{2\Phi_{max}} \right] \leq \frac{1}{\sqrt{n}}.$$

\square

SIS_{bs} : Lower Bound in Star-like Networks. Before we prove that the infection duration does not scale exponentially in star-like networks with high probability, we need the following lemma:

LEMMA 6. *In a clique of size n , for an arbitrary initial infected node, the second moment of extinction time for SIS_{bs} dynamics has an upper bound given by*

$$\mathbb{E}[\tau_e^2] \leq \frac{2(1 + \max(1, \Phi_{max}))^{2n}}{n\Phi_{max}^2}.$$

PROOF. We again consider the evolution of the total number of infected nodes with transition rates $q_{i,i+1} \leq (n-i)\Phi_{max}$ and $q_{i+1,i} = i+1$. The second moment for the first passage time to 0 starting from state 1 for a general birth and death process is derived in [30] and is given by:

$$\begin{aligned} \mathbb{E}[\tau_e^2] &= \frac{2}{q_{1,0}\pi_0} \sum_{i=1}^n \frac{1}{q_{i-1,i}\pi_{i-1}} \left(\sum_{k=i}^n \pi_k \right)^2 \\ &= \frac{2}{q_{1,0}\pi_0} \sum_{i=1}^n \frac{\pi_i}{q_{i,i-1}} \left(1 + \sum_{k=i}^{n-1} \prod_{l=i}^k \frac{q_{l,l+1}}{q_{l+1,l}} \right)^2. \end{aligned}$$

Using this equation, and $q_{i,i+1} \leq (n-i)\Phi_{max}$ and $q_{i+1,i} = i+1$, we have

$$\begin{aligned} \mathbb{E}[\tau_e^2] &\leq \frac{2}{n\Phi_{max}\pi_0} \sum_{i=1}^n \frac{\pi_0 \Phi_{max}^i \binom{n}{i}}{i} \left(\frac{\sum_{k=i}^n \Phi_{max}^k \binom{n}{k}}{\Phi_{max}^i \binom{n}{i}} \right)^2 \\ &\leq \begin{cases} \frac{2}{n\Phi_{max}^2} \frac{(1+\Phi_{max})^{2n}}{2^{2n}} & \text{if } \Phi_{max} > 1 \\ \frac{2}{n\Phi_{max}} & \text{if } \Phi_{max} \leq 1 \end{cases} \end{aligned}$$

\square

THEOREM 9. *Under the conditions of Theorem 6, $\log \tau_e = O(m \log n)$ with probability at least $1 - 2n^{-1} - n^{-\frac{m\Phi_{max}}{2e}}$ for large enough n .*

PROOF. As in the proof of Theorem 6, we consider cycles of epidemic evolution. We first bound the total number

of cycles and then give a bound for the duration of each cycle. As shown in Theorem 6, the total number of cycles is stochastically dominated by V , a geometric random variable with mean $\frac{1}{q}$. Let $K = Cn^{\frac{m}{\epsilon}\Phi_{max}} \log n$. Therefore, from Equation 3,

$$\mathbb{P}[V > K] = (1 - q)^K \leq \exp(-qK) \leq \frac{1}{n} \quad (4)$$

We next bound the duration of the first epoch in the cycle – the duration of the cycle in which at least one node in the core is infected. Since this time is stochastically dominated by the duration of infection in a clique of size m (size of the core), we have

$$\mathbb{E}[(T_{1,j} - T_{0,j})^2] \leq \frac{2(1 + \Phi_{max})^{2m}}{m\Phi_{max}^2}.$$

Using Markov inequality,

$$\begin{aligned} \mathbb{P}[T_{1,j} - T_{0,j} > 2n^{\frac{m}{\epsilon}\Phi_{max}}] &\leq \frac{\mathbb{E}[(T_{1,j} - T_{0,j})^2]}{4n^{\frac{2m}{\epsilon}\Phi_{max}}} \\ &\leq \frac{(1 + \Phi_{max})^{2m}}{2m\Phi_{max}^2 n^{\frac{2m}{\epsilon}\Phi_{max}}}. \end{aligned}$$

Using union bound, the total duration of these first epochs in K cycles can be bounded as follows:

$$\begin{aligned} \mathbb{P}\left[\sum_{j=1}^K T_{1,j} - T_{0,j} > 2Kn^{\frac{m}{\epsilon}\Phi_{max}}\right] &\leq K \frac{(1 + \Phi_{max})^{2m}}{2m\Phi_{max}^2 n^{\frac{2m}{\epsilon}\Phi_{max}}} \\ &\leq n^{-\frac{m}{2\epsilon}\Phi_{max}} \quad (5) \end{aligned}$$

for n large enough. Similarly, we upper bound the duration of the second epoch which is stochastically dominated by ζ_1 defined in Section 3.

$$\mathbb{P}[T_{0,j+1} - T_{1,j} > \frac{2C}{\epsilon} \log n] \leq ne^{-2C \log n}.$$

Therefore, the total duration of the second epochs in K cycles is bounded by,

$$\begin{aligned} \mathbb{P}\left[\sum_{j=1}^K T_{0,j+1} - T_{1,j} > \frac{2KC}{\epsilon} \log n\right] &\leq Kne^{-2C \log n} \\ &= \frac{C \log n}{n^C} < \frac{1}{n} \quad (6) \end{aligned}$$

for large enough n . Since $\tau_e = \sum_{j=1}^V T_{0,j+1} - T_{0,j} = \sum_{j=1}^V (T_{0,j+1} - T_{1,j}) + (T_{1,j} - T_{0,j})$, using union bound we have,

$$\begin{aligned} \mathbb{P}[\tau_e > 2K(n^{\frac{m}{\epsilon}\Phi_{max}} + C \log n)] &\leq \mathbb{P}[V > K] \\ &\quad + \mathbb{P}\left[\sum_{j=1}^K T_{1,j} - T_{0,j} > 2Kn^{\frac{m}{\epsilon}\Phi_{max}}\right] \\ &\quad + \mathbb{P}\left[\sum_{j=1}^K T_{0,j+1} - T_{1,j} > \frac{2KC}{\epsilon} \log n\right]. \end{aligned}$$

Finally, using Equations 4, 5, 6, we have for n large enough,

$$\mathbb{P}[\tau_e > 2K(n^{\frac{m}{\epsilon}\Phi_{max}} + \frac{C}{\epsilon} \log n)] \leq 2n^{-1} + n^{-\frac{m\Phi_{max}}{2\epsilon}}.$$

Since $\log(K(n^{\frac{m}{\epsilon}\Phi_{max}} + \frac{C}{\epsilon} \log n)) = O(m \log n)$, we have our result. \square